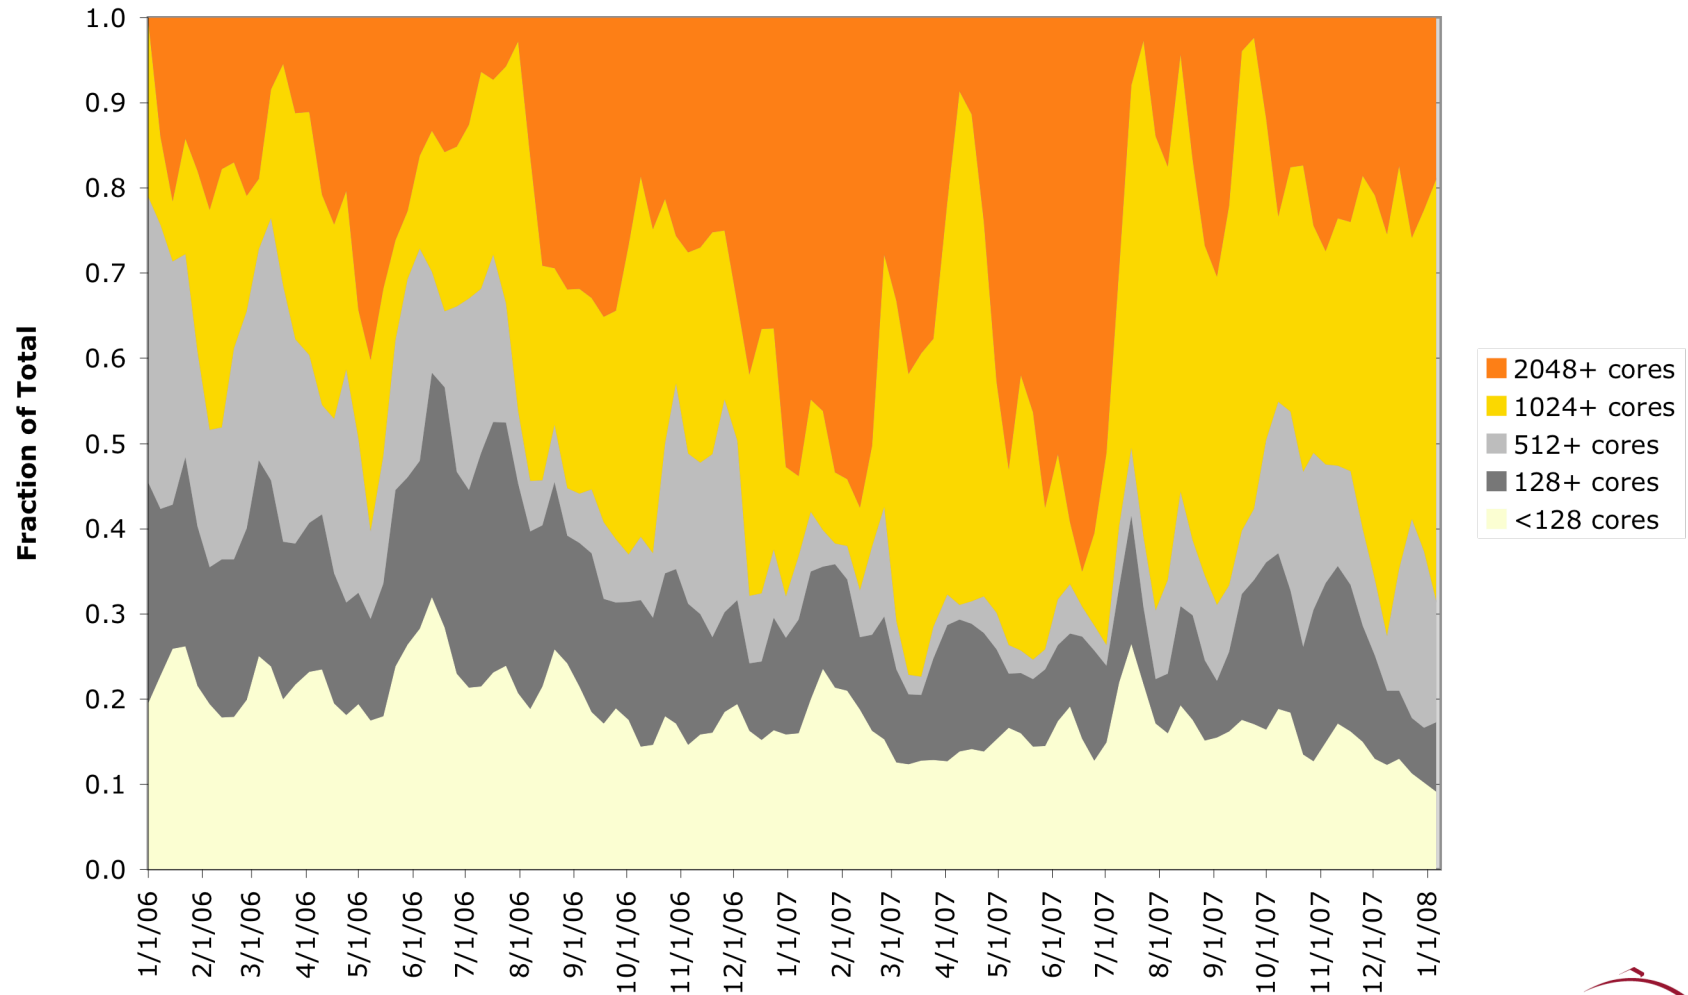


NERSC Machine Usage Characteristics

July 2008

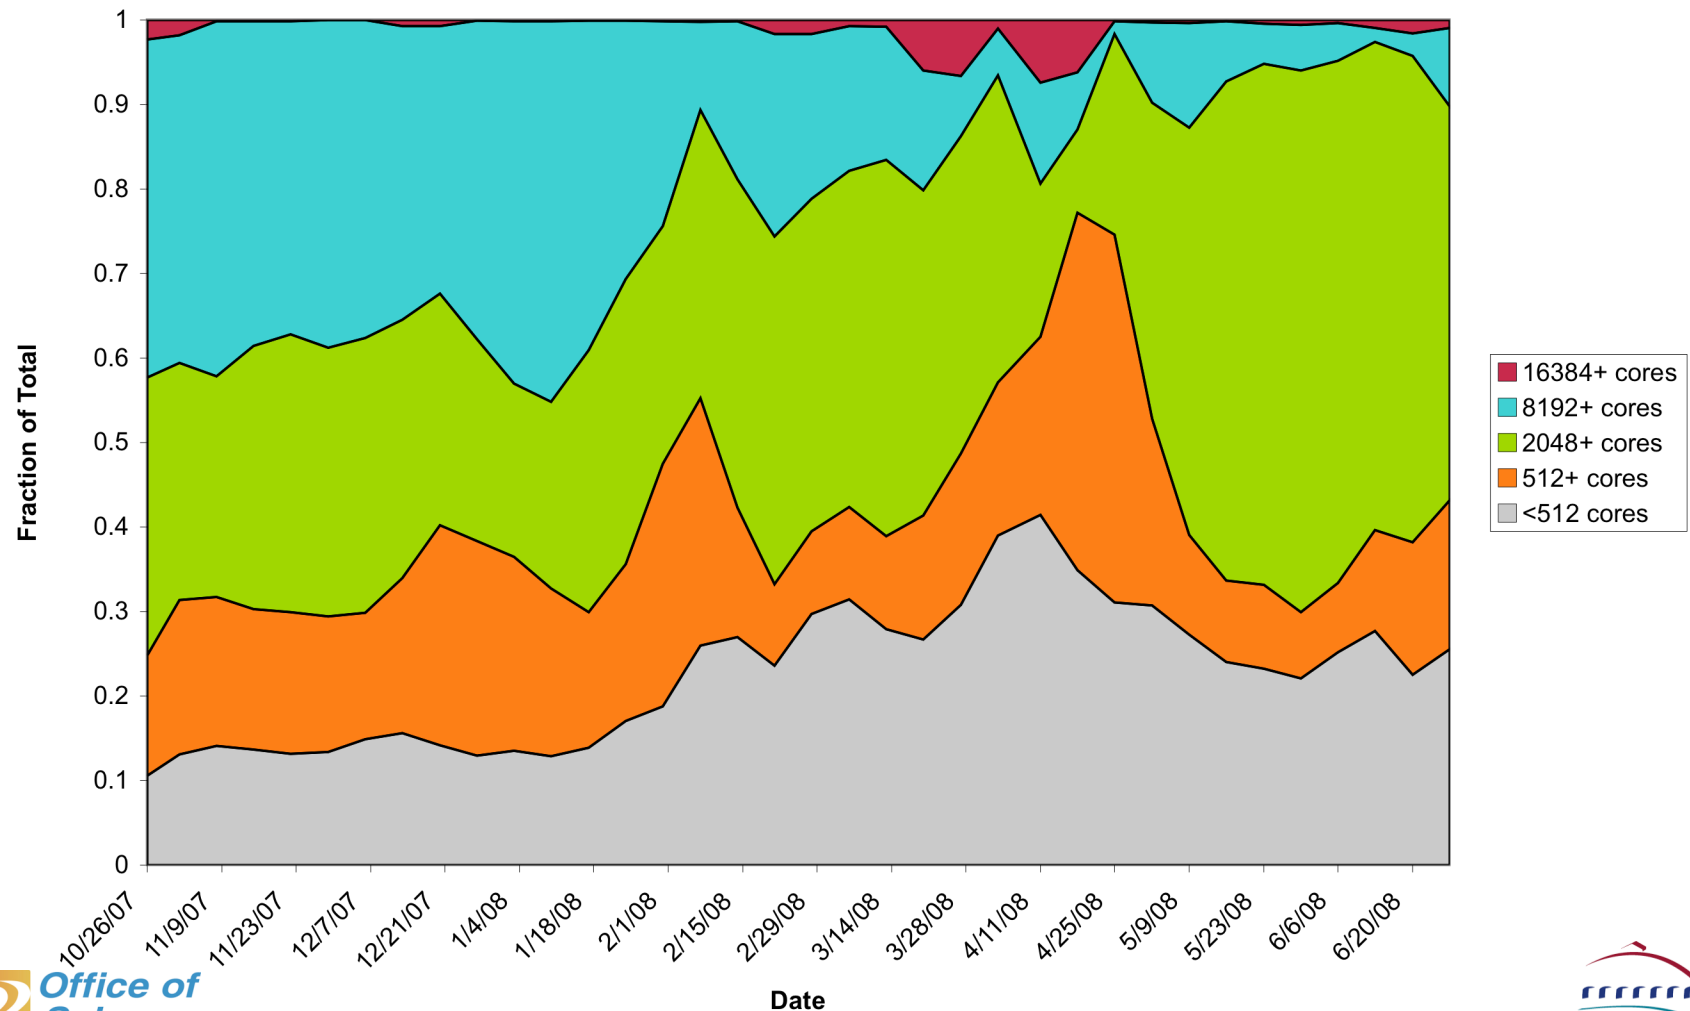
Computing at Scale

Seaborg Concurrency 06-07



Computing at Scale

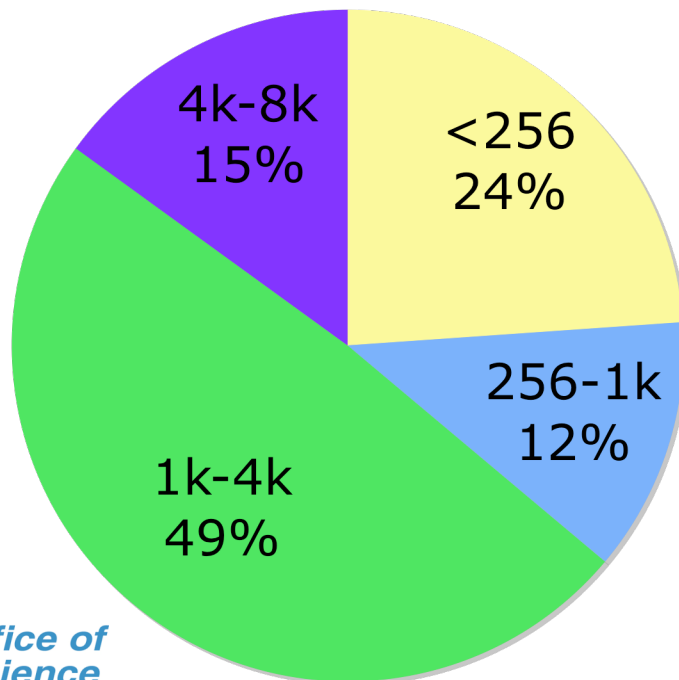
Franklin Concurrency



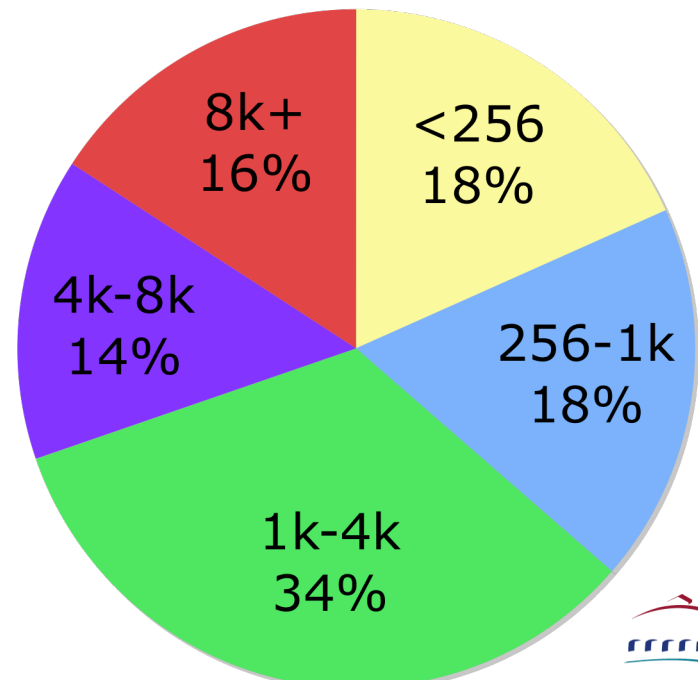
Computing at Scale

- Historically, almost 75% of the workload has run on >500 processors
- During the first six months of production, 50% of Franklin workload runs on >2000 processors

Job Concurrency on Seaborg (2007)

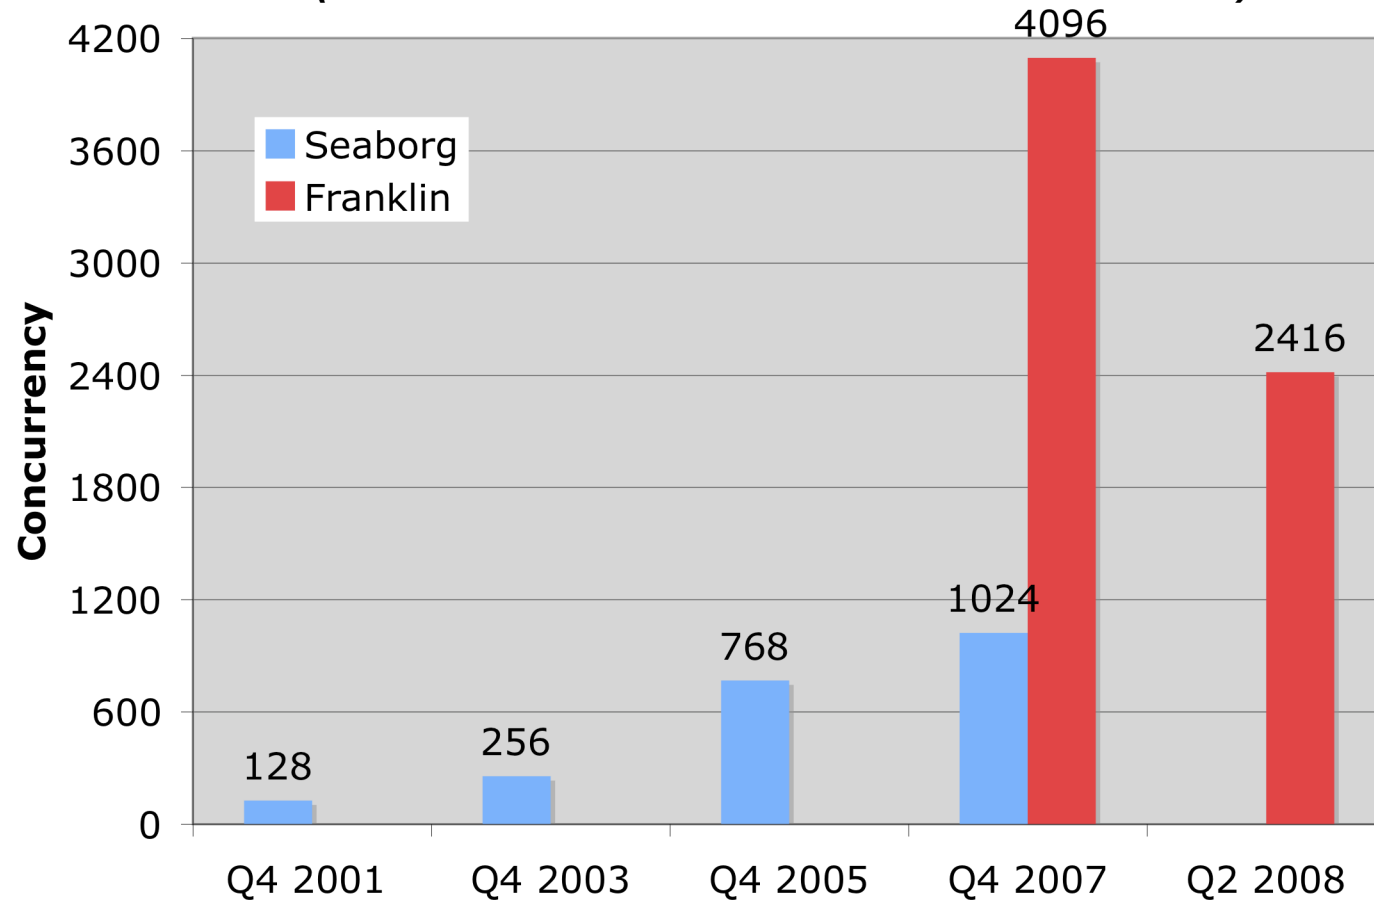


Job Concurrency on Franklin



Transitioning to Larger Concurrency

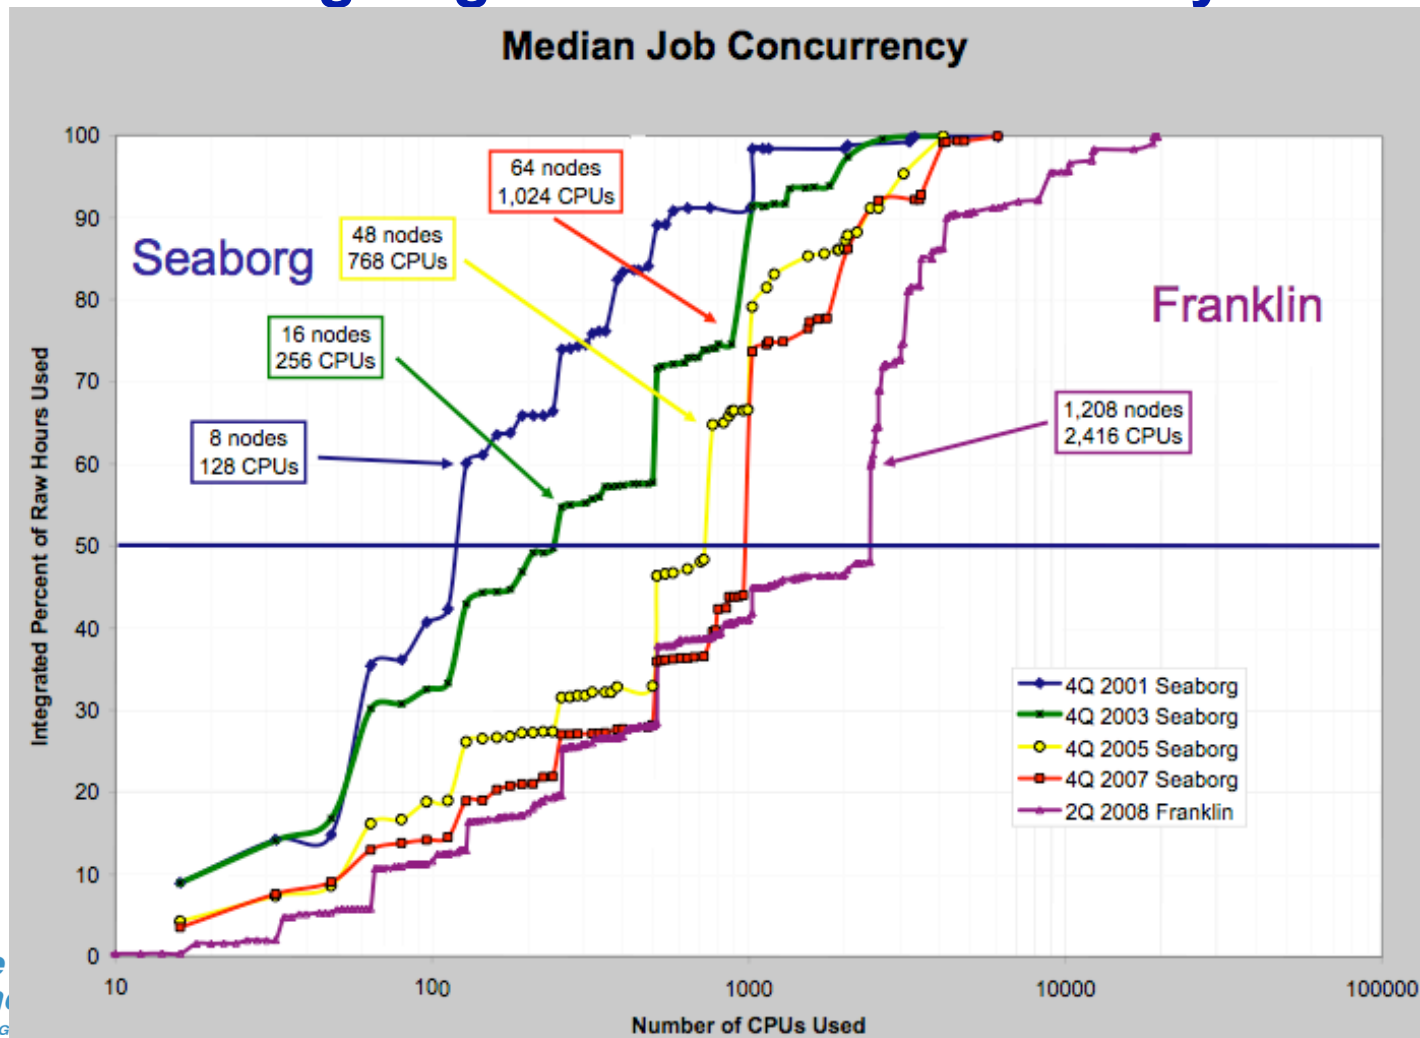
Median Job Concurrency
(Job Size where 50% of Raw Hours Used)



NERSC has enabled users to run at higher concurrencies

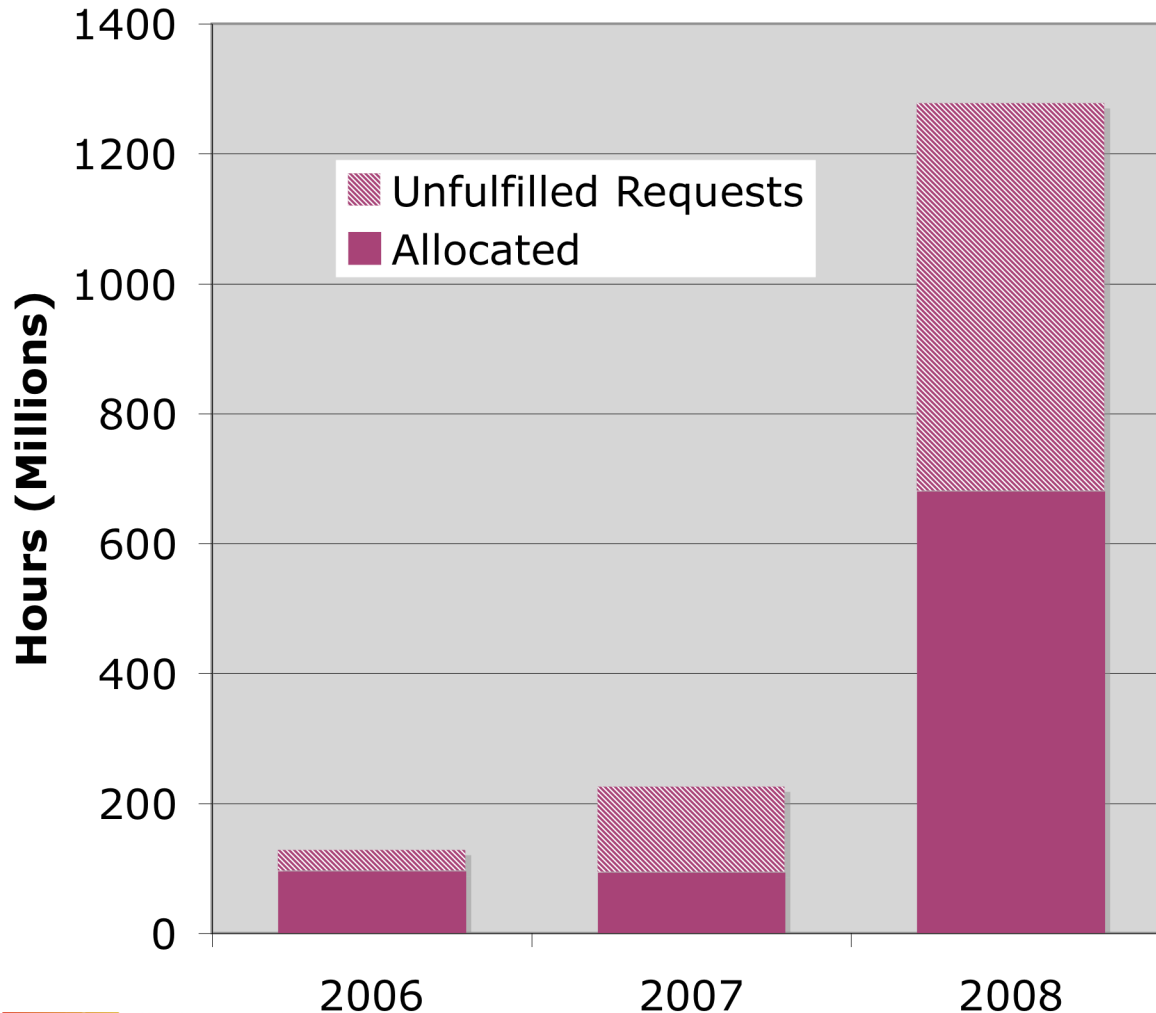
Transitioning to Larger Concurrency

DOE users have successfully transitioned to using larger concurrencies each year



Demand for More Cycles

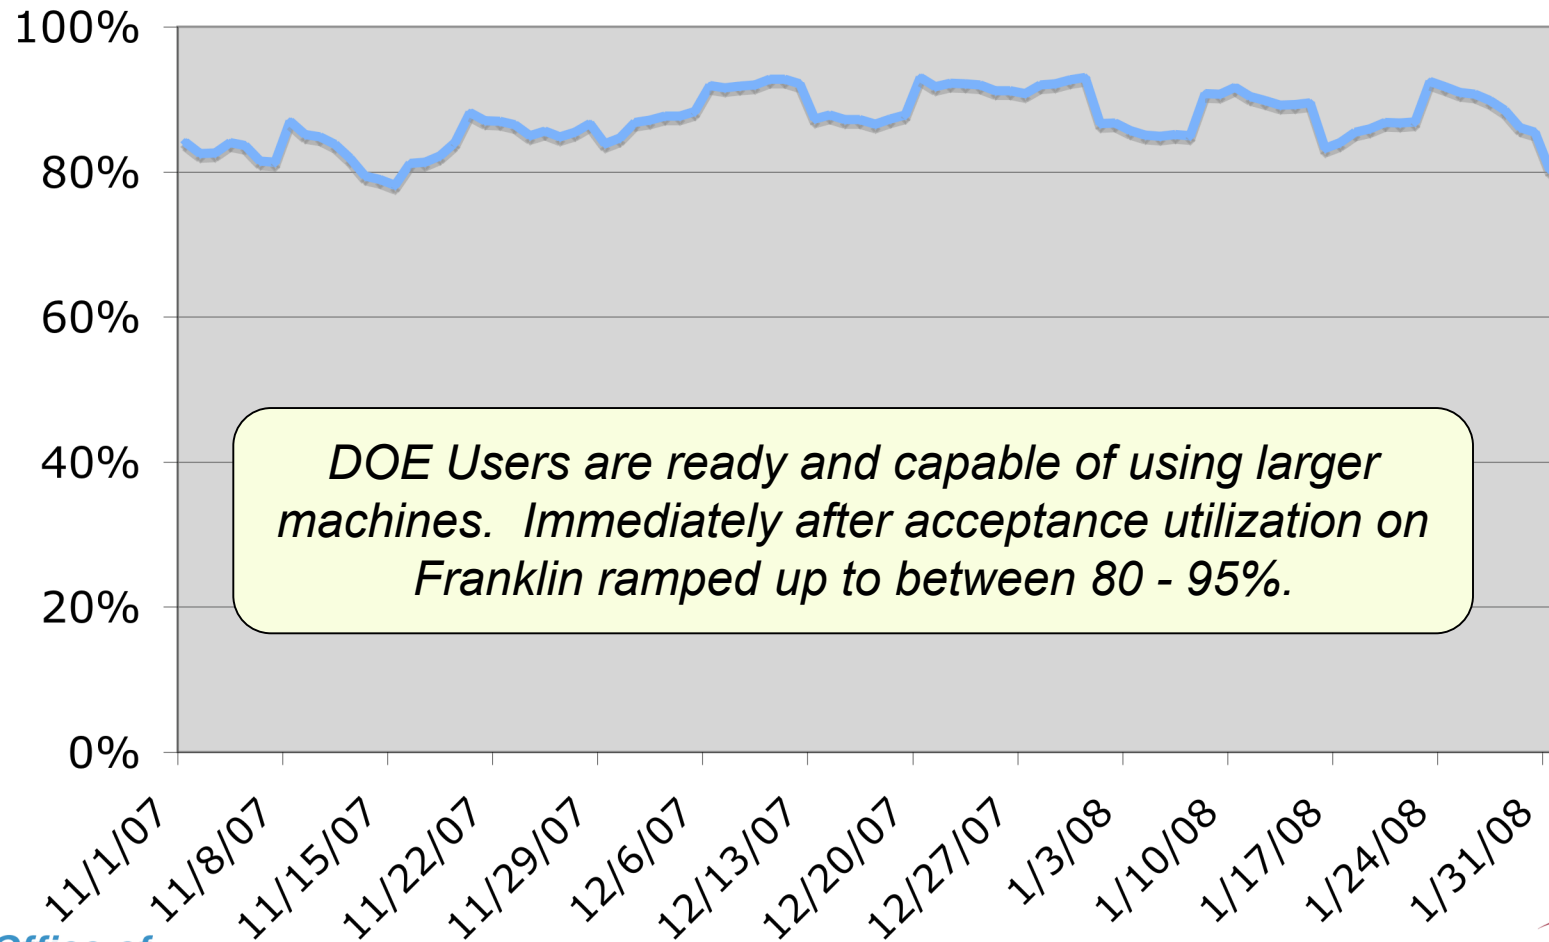
Compute Hours Requested vs Allocated



- *Each year DOE users request roughly twice as many hours as can be allocated*
- *Unfulfilled allocation requests amount to hundreds of millions of compute hours in 2008*

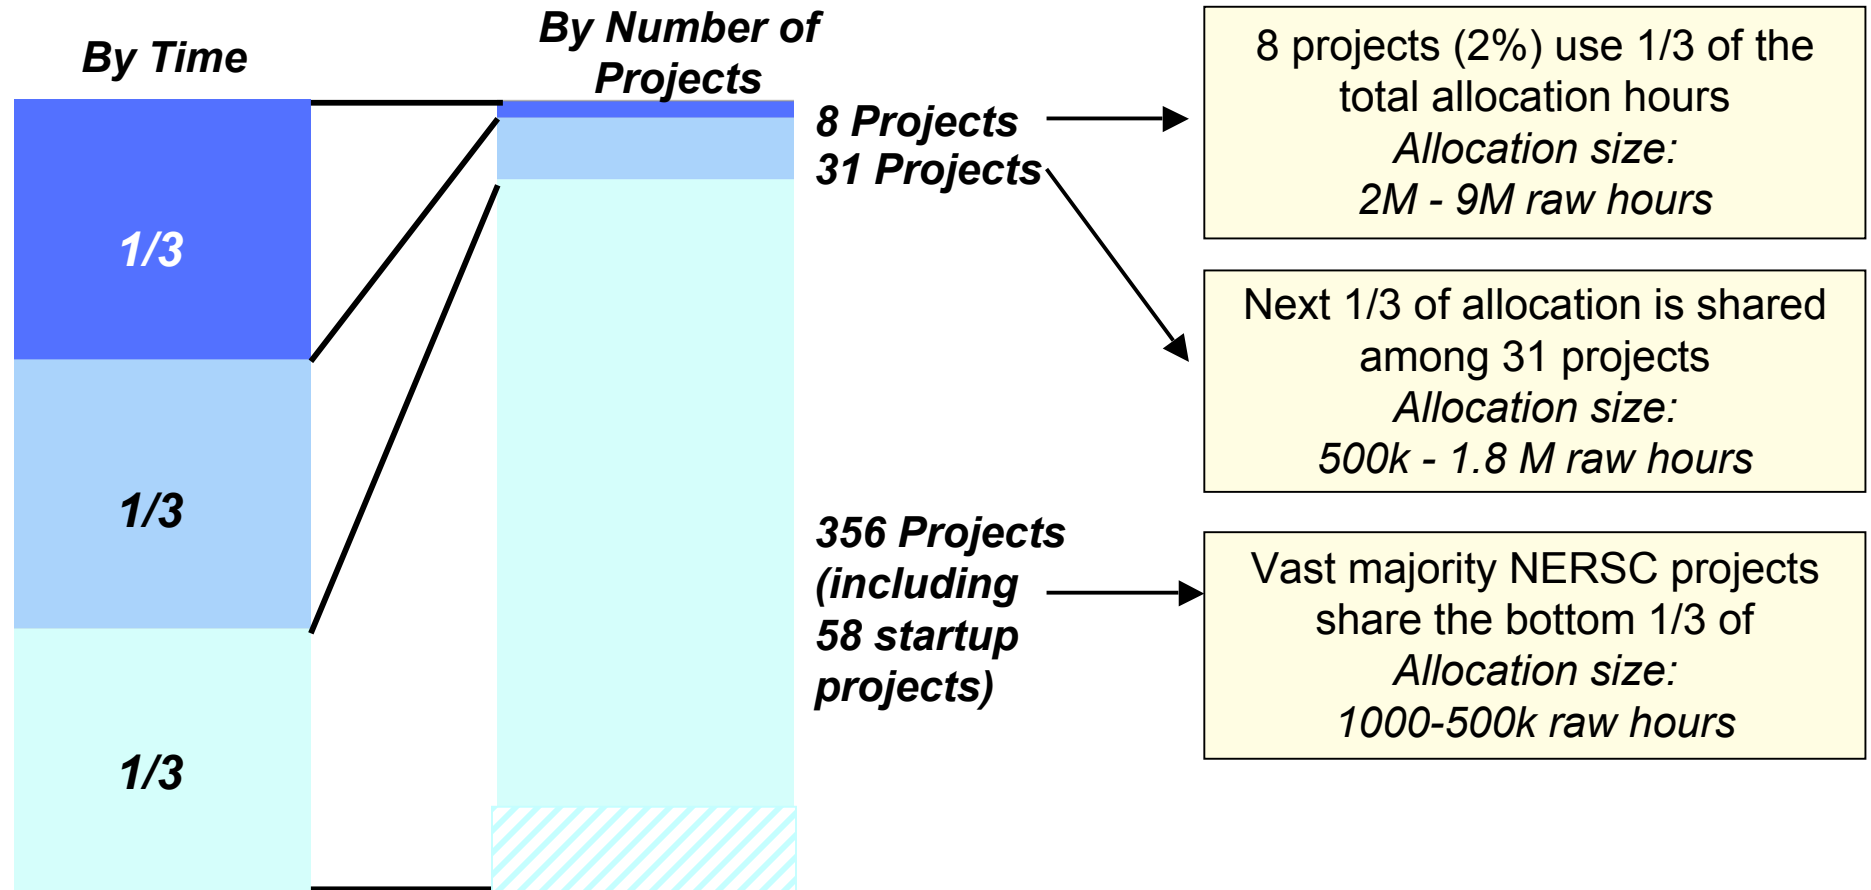
Fast Franklin Uptake

***Franklin Utilization for 3 Months After Acceptance
(7 day moving average)***



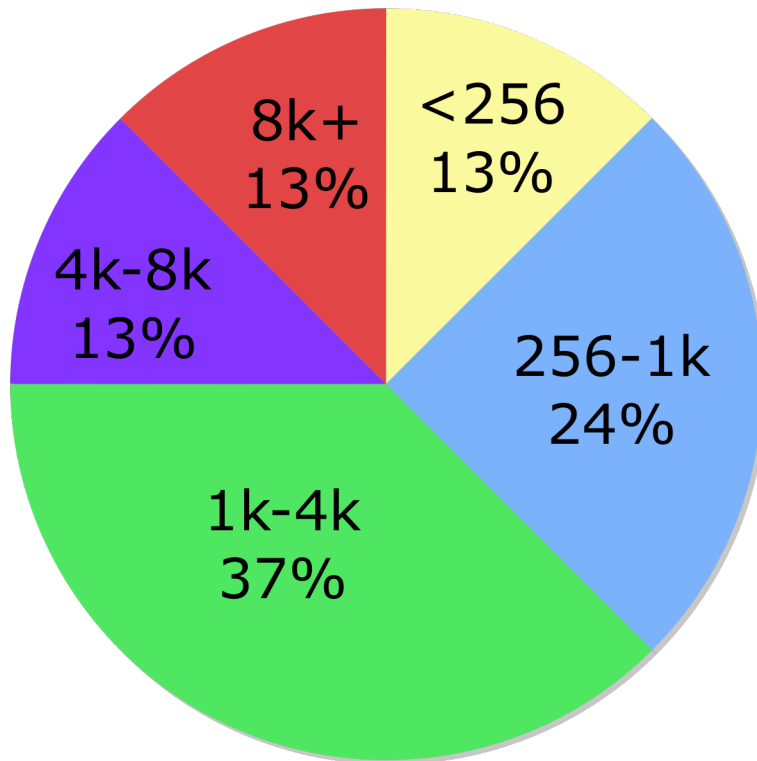
NERSC Allocation Breakdown

**Total NERSC allocation time
divided into thirds**

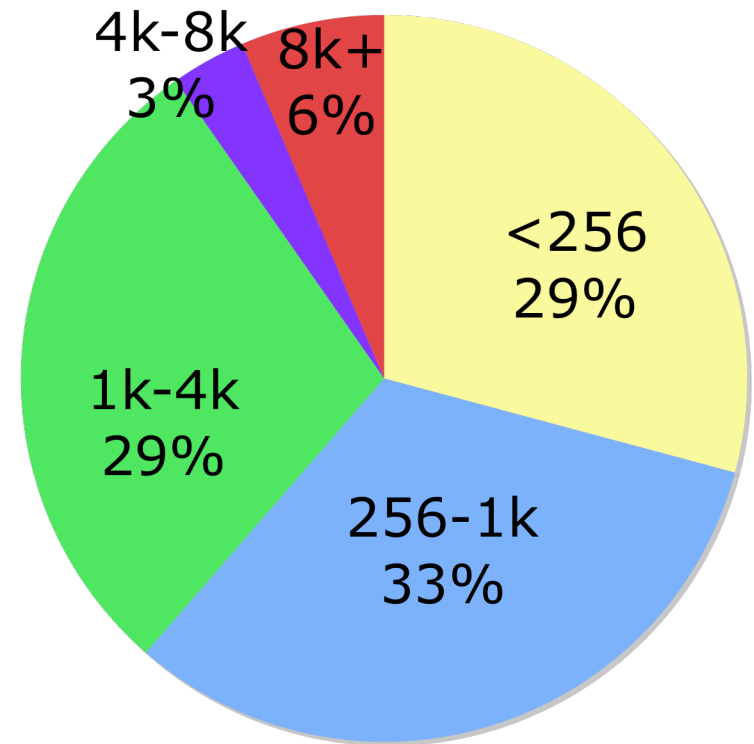


Concurrency on Franklin

Top Group of Users Most Frequent Concurrency



Middle Group of Users Most Frequent Concurrency



Both groups of users are capable of running high concurrency jobs ...